




ELŻBIETA KULIGOWSKA 
MACIEJ CZUCHRY
KRZYSZTOF CHYŻY 
TOMASZ KUNDERA 
CAROLE ROSKOWIŃSKI
MARTA A. DZIEŁAK

INSTALLATION AND TESTING THE LOFAR SOFTWARE ON ACC CYFRONET CLUSTER PROMETHEUS

Abstract *We present the actual status and the most important issues related to the installation of the data reduction LOFAR software on high power computer Prometheus located in ACC Cyfronet. We refer to the software itself as well its practical use cases in the context of the scientific tool and the detailed installation/testing methodology. We address most typical challenges and problems that occurred during our attempts to set up the complete and ready-to-use LOFAR environment (including not only programs, but also libraries, scripts and other additional tools) on non-standard (cluster-type) computing system. The result of these works is then briefly summarized. We also discuss the issues related to LOFAR documentation, maintenance, distribution and further development. Finally, we propose some future improvements.*

Keywords lofar, astronomy, imaging, radio astronomy, parallel calculations

Citation Computer Science 20(2) 2019: 165–177

1. Introduction

LOFAR (Low-Frequency Array¹) is the European network of low-frequency radio telescopes with the centre located in the Netherlands, completed in 2006–2012 by ASTRON (Netherlands Institute for Radio Astronomy). This large array uses the phenomenon of radio interferometry which results in providing the better resolved sky maps (images of selected celestial objects) due to the presence of many separate radio telescopes (or single antennas) electronically joined together to simulate one large antenna with the diameter comparable to the distance between the most remote elements of the system (the application of radio interferometry can theoretically produce images with the angular resolution of a system with an aperture equal to the separation between the component telescopes, provided that the gaps between detectors of the array are minimized [5]).

To achieve the expected long interferometry baselines and resolution, all the LOFAR components (antennas) are spread out over an area of more than 1000 km in diameter and localized in several European countries (Germany, France, United Kingdom, Sweden, Poland and Italy), with the computational and logistic center located near the Dutch town Groningen. To outline the complexity of this huge device: LOFAR is made up of about 25000 simple antennas grouped in fields (stations), with each single station consisting of two sets of 96 antennas (working with both lower and higher observing frequencies). The obtained data are radio fluxes of celestial bodies (typically: radio galaxies, AGNs, Sun etc.) collected at radio waves from the range of 10 to 240 MHz, with the highest so far achieved resolution. LOFAR is the most sensitive radio observatory at its low observing frequencies and also constitutes the test and demonstration of the technology intended to be used in the new generation Square Kilometer Array interferometric radio telescope [2].

From the scientific point of view, LOFAR is particularly useful in the innovative research projects related to the objects of the distant Universe created shortly after the Big Bang: most remote galaxies and the earliest structures (clusters of galaxies and galaxies themselves, including radio galaxies at high redshifts). Its probably most exciting and technically challenging applications is studying the redshifted 21 cm line emission from the Epoch of Reionization – the period after recombination when the Universe turned neutral. However, the expected signal (redshift range from $z = 11.4$ (115 MHz) up to $z = 6$ (200 MHz) is weak and needs to be disentangled from the much stronger foreground (ie. galactic) emission.

The question of installing the data reduction LOFAR software on superclusters Zeus and Prometheus (located in the Academic Computer Centre (ACC) CYFRONET in Cracow, Poland) has been first taken into consideration under the PLGrid Plus project (2011–2015) [1] and then continued until the early months of 2017. The main argument for installing this software on high-power computes concerns the very large data streams typically produced in LOFAR astronomical obser-

¹www.lofar.org

vements. Such amounts of data entails the use of large computing resources that are often inaccessible in the case of working stations or personal computers. In this type of scientific research the superiority of high-power computers over classical desktop computers should be also visible in increasing computational speeds and the possibility to use so-called parallel calculations, presumably already implemented in the LOFAR packages and tools. Having the LOFAR environment installed on a high-power cluster computer, one can expect: (a) facilitating in the storage of LOFAR data files on disk resources (b) significant acceleration of data processing due to the dramatically increased memory resources (c) further acceleration of the calculation due to the parallelization and multi-threaded processors. All of the above summed together should result in speeding up the entire process of producing the detailed radio maps of the celestial objects, with particular emphasis on the extended radio galaxies, usually demanding also multifrequency observations (see Figs. 1, 2).

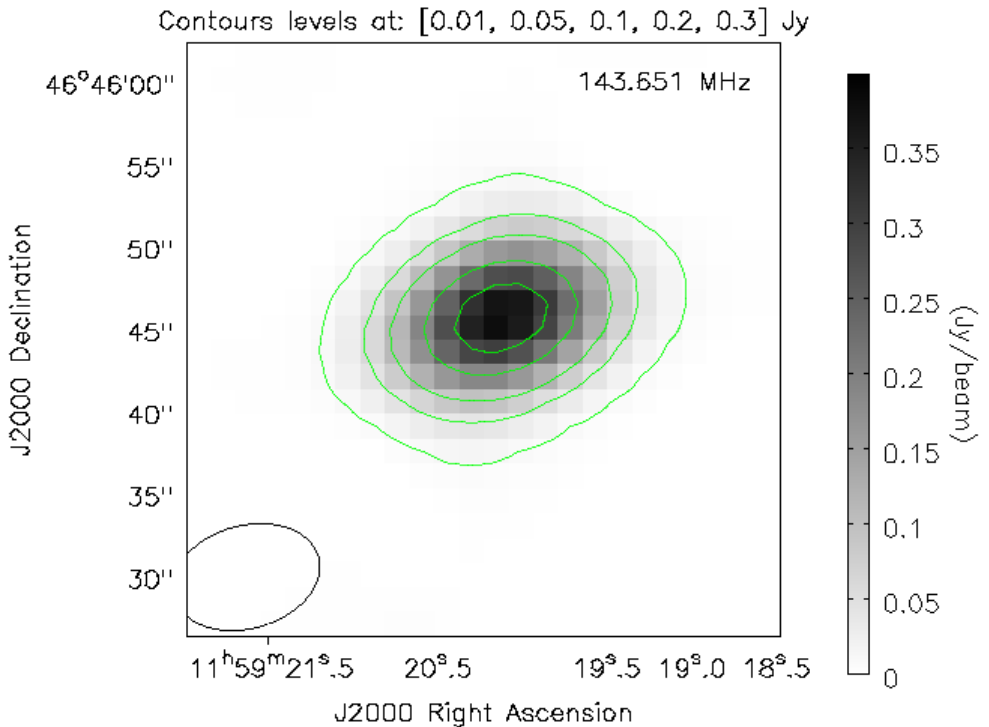


Figure 1. 1159+4645 (B3 1156+470): young compact radio galaxy from the family of CSS/GPS OBJECTS observed by LOFAR telescope at 143.651 MHz. The maps were created with the use of all available subbands (up to 240 SBs and 8 hours of observation) and with FACTOR




Figure 2. Preliminary release LOFAR LoTSS [6] mosaic covering the P15 field with the 1159+4645 radio source. LoTSS (LOFAR Two-metre Sky Survey) is a deep 120–168 MHz frequency imaging survey conducted to cover the entire Northern sky. Each of its 3170 pointings will be observed for 8 hours which at most declinations is sufficient to produce ~ 5 arcsec resolution images with a sensitivity of ~ 0.1 mJy/beam and accomplish the main scientific aims of this survey: exploring the formation and evolution of massive black holes galaxies clusters of galaxies and other large-scale structures in the Universe. Figure presents the LoTSS zoom in/zoom out with the exact position of 1159+4645

To test the above hypothesis, we conducted the complete installation of the LOFAR computing environment on ACC Cyfronet computing cluster Prometheus². We describe this process in Section 3. We then tested the compatibility of the scientific results obtained with LOFAR software installed on Prometheus with the ones provided by traditional computing systems, including linux-based desktops and high-end workstations.

Finally, we checked whether the average data reduction speeds are actually significantly higher in the case of the supercluster Prometheus. We summarize this step in Section 4.

2. LOFAR: software usage opportunities

2.1. Data processing at the dutch operational center

LOFAR Data collection and early processing is (mostly³) performed at the Dutch University of Groningen and with the use of the new (2013) COBALT (CORelator and Beamforming Application platform for the Lofar Telescope) CPU-GPU based system as the central correlator and beamforming platform for the International LOFAR

²<http://www.cyfronet.krakow.pl/computers/15226,artykul,prometheus.html>

³There are also local national data centers, for example polish POLFAR

Telescope⁴. The individual array elements (stations) are connected by a fast fiber optic (backbone) network.

In theory, data post-processing can be carried out with both high performance user computers and dedicated computing clusters. However, taking into account the necessity to achieve the satisfactory computation time (for large datasets extending up even to several months at “local” computers), modern supercomputers with limited access are preferable. In the Netherlands, all the data preprocessed by the Correlator are subsequently placed to a cluster of computer nodes CEP4 designed for the exclusive use to process the data through the standard data reduction pipelines. The separate CEP3 cluster is also available for manually analyse of the incoming data⁵.

In the case of CEP4, the total storage provides a 2PB LustreFS global filesystem through the InfiniBand network to all nodes and the users processing jobs are run under SLURM batch scheduling system. Both CEP 3 and 4 are generally access-limited to the users accepted by Radio Observatory’s Science Operations and Support group (SOS). Access is granted only at the discretion of the SOS group and after submitting a formal request (most probably due to the huge amount of resources consumed by a single user). Additionally, CEP3 awarded users are usually able to access it for a very limited time (8 weeks by default). For the above reasons the constant use of CEP3/CEP4 resources by international users does not seem to be optimal. In addition, there is also the problem of sending large amounts of data between the described dutch computing center and local (here: polish) LOFAR observing stations. Thus it seems much more reasonable to use rather the local computing resources instead of the central ones.

2.2. LOFAR software – the general characteristics

LOFAR software is composed of several modules executing the individual steps of data reduction: data preprocessing (flagging, averaging, demixing), data calibration (using the known calibrator which is usually the source with well confirmed radio flux), and the final procedures for creating two dimensional radio maps (the intensity of radio flux at the given range of frequencies and at the given sky area limited by defined sky coordinates: right ascension and declination) resulting from previously processed UV plane coverage (the visibility function in the Image Synthesis theory [7]). In a classic approach, every data reduction step needs to be followed by careful data inspection before proceeding to the next one (ie. to evaluate the quality of the solutions achieved

⁴COBALT system consists of 8 production nodes and 1 test node, each consisting of 2 CPUs (Intel Dual Xeon E5) and 2 GPUs (NVIDIA K10) housed in a DELL T620 box and connected by an FDR Infiniband Switching network. This gives a balanced system where each CPU connects to one GPU, one Infiniband port and two Ethernet ports and both CPUs within one node are also connected. The cooling of the GPU cards is achieved due to special air ducts constructed by the CIT Groningen en the ASTRON Mechanical Department. Full system is certified by DELL.

⁵An extensive description of CEP4 and CEP3 can be found at <https://www.astron.nl/radio-observatory/observing-capabilities/depth-technical-information/cep-and-lta-computing-facilitie>.

so far). However, in practice it involves the necessity to ensure the instant access to the graphical preview of the current results – in some cases also during the calculations themselves. This requirement of interactivity is a new challenge, especially when compared to more traditional cluster computation usually using the input set of parameters and data only to be processed.

The above-described problem can be partially solved and the calculation performed in the frame of data reduction may be automatized. Over the past years, LOFAR scientists has adopted a special processing model in which final data products for routine observations are formed in automated mode, with user interaction needed only at crucial moments (so-called FACTOR [3]). However, for those seeking to implement these advanced solutions for independent clusters the major challenge is the fact that LOFAR architecture is still being defined and developed with frequent changes. It implies that the complete and accurate documentation referring to its installation is very difficult to provide and maintain, shifting some part of the installation–related works to scientific LOFAR users.

2.3. Using the LOFAR package on local computers

To make comparisons possible and to estimate the expected improvements in speed of calculation we have created two installation as a reference. First of them was maintained as a very basic test of installability and library dependencies checker. It consisted of standard desktops under control of the Linux system starting from the old Intel P4 platform up to newer based on Intel i5 processors with 3.3 GHz clock, 8 GB RAM, and medium speed rotational disks. The OS used was Scientific Linux (Red Hat Enterprise Linux derivative created in CERN) versions 5 and 6. Because these systems were used primarily as first sandboxes for LOFAR installation, they were not highly optimized though they have been used as the reference point “zero”. Obviously test runs were possible on them, but effective calculation times for data reduction processes were not practical with typical length of weeks.

To measure actual computational time on workstation class computers and to create a more accurate reference point a few more powerfull machines have been used. They were two types of workstations, built on ASUS Z9PA boards with double Intel Xeon E5420 2.5 GHz and double Intel Xeon E5-2640 2 GHz CPUs. They have been equipped with 16 GB or 128 GB RAM accordingly with different memory layouts (less number of bigger 32 GB chips, or more numerous 16 GB ones). It has been found that memory layouts does not produce measurable differences. Each workstation had its own hard drives set based on server class Western Digital Red series disks with 3 TB of capacity. The operating system used was Linux Ubuntu 14.04 LTS, as most of the LOFAR software was produced in similar environment⁶. Both hardware and software were as far as possible optimized for calculation speed. As big amount of RAM has been found crucial for some operations hence in that cases workstations

⁶It has been decided lately by software maintainers to switch the base environment to RHEL derivatives namely CentOS.

with bigger RAM amount have been used, while others operations were performed on all the computers of the infrastructure.

3. LOFAR installation on ACC Cyfronet AGH cluster Prometheus

To ensure satisfactory calculation speed and suitably large disk resources, an effort was made to install the complete LOFAR environment (including the still-developed FACTOR module) on ACC Cyfronet AGH cluster Prometheus (configuration HP Apollo 8000) with the following resources: 279 TB or RAM memory, 10 PB of total disc storage space, and computing power equal to 2399 Tflops. The installation of most of the LOFAR components on Prometheus was performed based on the procedures from the LOFAR wiki pages⁷, providing the instructions for Ubuntu 14.04 and CentOS 6.5. Installation was carried out in a hybrid manner. First, it was done using the popular Intel compilers (icc, icpc) for C, C++ and Fortran-written components. However, due to the incompatibility of this components with some parts of the additional LOFAR code, GNU GCC (gcc and g++) compilers had to be also used for C and C++-written components, while the Fortran components compiler was Intel ifort (due to its compatibility with some libraries previously installed from the sources). Further, some parts of the code (with MPI support) were build with Intel MPI. To achieve the proper functioning of the software on the cluster it was also necessary to modify the selected Python libraries and adapt them for the SLURM-based environment (typically not present on linux-based desktops).

After this primary installation, the default and common modules and tools used for data reduction were tested. These components include: CASA/Casacore, Python libraries, BBS-reducer used for data calibration, NDPPP (Default Pre-Processing Pipeline) for data flagging, and final mapping tool AWImager. We concluded that the proper installation of these components also demanded setting up some user-defined paths (not always described in official documentations) and adding a number of additional libraries and dependencies (such as log4cplus, WCSlib, Boost, AOFlogger, WSClean, PySide or PyQt4). After these first tests we also realized that some additional scripts and files have to be added to linux PATH variable to obtain the full software's productivity. These are:

- addqueue.sh,
- addtopic.sh,
- asciistats.py,
- calibrate,
- calibrate-part,
- calibrate-stand-alone,
- calibration-pipeline.py,
- casaplotms,

⁷<https://www.astron.nl/lofarwiki/doku.php?id=start>

- cleanupq.sh,
- delqueue.sh,
- deltopic.sh,
- dpuxml-interface.py,
- fed.sh gsm.py,
- genericpipeline.py,
- imaging-pipeline.py,
- listqueues.sh,
- loader.py,
- long-baseline-pipeline.py,
- msss-calibrator-pipeline.py,
- msss-imager-pipeline.py,
- msss-target-pipeline.py,
- pcombine.py pipeline-job.py,
- preprocessing-pipeline.py,
- pulsar-pipeline.py,
- qls.sh,
- QtoPipeline.py,
- selfcal-imager-pipeline.py,
- startPython.sh,
- startPythonVersion.sh,
- statsplot.py,
- stopPython.sh.

Finally, it is worth noting that proper LOFAR environment demands setting up the local GSM database with MonetDB installation⁸.

4. Further tests of the LOFAR ACC Cyfronet installation

Prometheus operating system is CentOS 7. Thus most of the scientific packages (e.g. Boost) needed to be installed manually instead of using repositories. It implies that all the paths and variables had to be set correctly (to create ready environment for LOFAR packages) during the installation process itself. To obtain such environment for building and running this software, we decide to use Lmod environmental modules system [4]. This solution enables to load (for example) complete LOFAR environment with all its dependencies with one command (*module load...*).

It should be also emphasized that in the early years of LOFAR the process of data reduction was highly interactive: it required frequent access to the lower-level data products and its careful evaluation. Currently the work is continuing on to streamline these processes that are supposed to be controlled by the human

⁸<https://www.monetdb.org/Documentation/Guide/Installation>

user at a lesser extent. For example, FACTOR module is a tool for automatized producing high-resolution images from LOFAR data [7]. We attempted to install also this component. FACTOR installation on Promethues was rather standard though it was necessary to add the full paths to dependencies during the setup process. According to information given by GitHub, FACTOR is “expected to run well” on both local user machines and multiple node clusters. However, despite the fact that the installation of Lofar FACTOR on Prometheus seems to be complete, its detailed test shows clearly that the entire processing pipeline is not working properly at some point. We conclude that FACTOR probably needs more sophisticated technical and scientific tests aimed to find out the source of this problem. It would be also useful to estimate the rates of FACTOR execution for the various reduction steps in comparison to similar rates obtained at local machines.

Scientific test of the general Lofar environment also raised the interesting conclusion: some of the LOFAR pipeline components are not uniform in the sense of parallel computing capabilities. The final data reduction step (realized with AWImager) seem to be not parallelized and running it on the cluster system does not bring significant time gain, particularly when compared to the BBS calibration module. However, it is also noteworthy that none of the LOFAR modules directly uses parallelization based on processes/threads. Instead, some of them are loaded with genericpipeline.py script that runs many instances of the given application simultaneously (each of them using one processor, and thus efficiently accelerating the overall calculations). It is not in our competence to evaluate this solution, but the full parallisation of all LOFAR components may be also highly useful in the case of future data processing projects, demanding even more computational resources [2].

4.1. Testing examples and use cases

All the tests were carried out using three examples of the previously reduced exemplary data (in order to easily compare scientific results) and covered: validation of the individual libraries and scripts, the quality of graphical display for input and output data, the usefulness of widely used (though partially not documented) Python scripts for data management, and the correctness of the execution of the core LOFAR modules as well as correctness of these modules (their resulting scientific results). We conclude that the Prometheus Lofar environment works properly, with the exception of the FACTOR module which still demands further checks and tests.

4.2. Comparison to the previous solutions

We also checked whether the Lofar installation Lofar on Prometheus cluster actually results in a noticeable time gain. Here we use the HBA data (MSSS – Multifrequency Snapshot Sky Survey project) observed on 29 March 2013 from 14:01:48.0 to 14:08:46.6 UTC and centered at 01 02 21.73 +31 27 36.00 (J2000), with the central frequency of 134.3735 MHz.

The rates for the exemplary calculation steps run on Prometheus compared to the same processes run on the local workstation (with 32 cores of Intel(R) Xeon 2.00GHz processors and 128 GB disc space) are presented in Table 1.

Table 1

Prometheus versus local machines: duration of the same exemplary LOFAR processes

	Prometheus	Workstation (OA UJ)	Workstation (PCSS)
BB-reducer	1h 10m	2h 12m	1h 27m
MWImager	1h 46m	2h 2m	1h 52m
fixbug.py (script)	1m 8s	4m 41s	1m 54s

It needs to be emphasized that the below rates apply to the relatively small amount of data to process, therefore are not showing the real time gain (it is expected that acceleration of computing time will be evident for more sophisticated initial data, especially composed of many subbands). The most conclusive result would be in this case comparing the similar rates for executing the FACTOR module, which is still not operating at this time.

4.3. Further impact

The overall experience gained during our attempts to install the LOFAR environment on ACC Cyfronet cluster Prometheus was directly used to perform the faster and successful LOFAR installation on the computational virtual machine located in PCSS (Poznań Supercomputing and Networking Center⁹). This workstation is characterized by the following parameters: 16 vCPU Intel(R) Xeon(R) CPU E5-2620 v2 @ 2.10GHz, 16 GB RAM and 250 GB of the total disc storage. In this case, the installation and testing process was finalized in about a month.

5. Conclusion and (suggested) future works

LOFAR installation on ACC Cyfronet cluster Prometheus is now partially useful from a scientific point of view. Its further tests should focus on estimating the exact speeds for the individual reduction steps of the FACTOR module, as well as completing the satisfactory installation of this module to ensure its full utility. We also conclude that the parallelization of all the LOFAR components (both used in FACTOR and in the “manual” mode) would significantly increase these speeds. In particular, it has been noted that some components tend to work much faster than the same components loaded on the local workstations, enabling more effective calculations, but this is not a general rule for every single LOFAR component tested on Prometheus.

⁹<http://www.man.poznan.pl/online/en/>

The main inconvenience encountered during the installation process was due to the fact that LOFAR itself is a living (continuously developed and upgraded) environment. Some of its parts are also very poorly documented, or demand a certain scientific knowledge to be not only used, but also set up properly. Software installation and maintenance may be also problematic due to the lack of the sufficient support for different architectures and systems (with emphasis on clusters)¹⁰. With respect to all the difficulties faced during setting up the LOFAR environment on Prometheus, the following proposals and ideas for improving the process of the general LOFAR installation at high performance computers may be suggested to its developers to be considered for the future LOFAR (or LOFAR-like) releases:

1. Preparing the proper configuration files for software installations executed with various Intel compilers (actually GCC i Clang are officially supported).
2. Completing the satisfactory and frequently updated documentation for LOFAR software, including FACTOR and other new modules as well as the instructions for installing/upgrading the individual additional components, databases and libraries that are not included in the original installation packages.
3. Implementing software support for accelerative libraries such as Intel MKL Xeon Phi or NVIDIA CUDA (GPGPU). For example, CUDA is a parallel computing platform and programming model that enables dramatic increases in computing performance by using the power of the graphics processing unit (GPU). It should be then obviously useful in the case of dedicated scientific environments like LOFAR.
4. (Time-consuming but probably possibly to done only once): Preparation of the ready binary packages for various popular Linux distribution (RHEL, Ubuntu, Debian, CentOS etc.) for both high-performance computers and high-end workstations. This approach should make using LOFAR easier and more accessible to young scientists and students.
5. Distributing LOFAR software through CVMFS – caching http based read-only file system optimized for delivering experimental software to (virtual) machines, originally developed as part of the CernVM project¹¹ (possibly in connection with point 2).
6. Global database for computational results with access restricted to authorized users (to be considered by scientists).
7. Finally, regarding the presence of not parallelized or not fully parallelized parts of the LOFAR pipeline, it could be very interesting to consider its possible parallelization in order to accelerate the most time-consuming calculations.

¹⁰Maintainers of the core LOFAR software have lately decided to switch to CentOS as their basic software environment. We hope this will make it much easier to install new versions of LOFAR software on Prometheus due to similar base OSes.

¹¹<https://github.com/cvmfs>

Acknowledgements

The research presented in this paper was partially supported by the Polish National Science Centre grant No. 2013/09/B/ST9/00599 and the Domain-oriented services and resources of Polish Infrastructure for Supporting Computational Science in the European Research Space – PLGrid Plus (2011–2015).

References

- [1] Bubak M., Kitowski J., Wiatr K. (eds.): *eScience on Distributed Computing Infrastructure: Achievements of PLGrid Plus Domain-Specific Services and Tools*, Springer, 2014.
- [2] Cordes J.: The Square Kilometer Array: Key Science and Technology Development. In: Kassim M.J.N., Perez M., Henning P. (eds.), *From Clark Lake to the Long Wavelength Array*, Bill Erickson's Radio Science ASP Conference Series, vol. 345, Proceedings of the Conference held 8–11 September, 2004 in Santa Fe, New Mexico, USA, p. 461, 2005.
- [3] van Haarlem M., et al.: LOFAR: The LOw-Frequency Array, *Astronomy & Astrophysics*, vol. 556, p. A2, 2013.
- [4] Mechev A., et al.: An Automated Scalable Framework for Distributing Radio Astronomy Processing Across Clusters and Clouds. In: *International Symposium on Grids & Clouds (ISGC 2017), Taipei, Taiwan, 2017*.
- [5] Röttgering H.: LOFAR, a new low frequency radio telescope, *New Astronomy Reviews*, vol. 47(4–5), pp. 405–409, 2003. [http://dx.doi.org/https://doi.org/10.1016/S1387-6473\(03\)00057-5](http://dx.doi.org/https://doi.org/10.1016/S1387-6473(03)00057-5).
- [6] Shimwell T., et al.: The LOFAR Two-metre Sky Survey. I. Survey description and preliminary data release, *Astronomy & Astrophysics*, vol. 598, p. A104, 2017.
- [7] van Weeren R.J., et al.: LOFAR Facet Calibration, *The Astrophysical Journal Supplement Series*, vol. 223(1), p. 2, 2016. <http://dx.doi.org/10.3847/0067-0049/223/1/2>.


Affiliations

Elżbieta Kuligowska 

Astronomical Observatory of Jagiellonian University, ul. Orla 171, 30-244 Kraków, Poland, elzbieta@oa.uj.edu.pl, ORCID ID: <https://orcid.org/0000-0002-8407-8280>

Maciej Czuchry

Academic Computer Centre CYFRONET AGH, Nawojki 11, 30-950 Kraków, Poland, m.czuchry@cyfronet.pl

Krzysztof Chyży 

Astronomical Observatory of Jagiellonian University, ul. Orla 171, 30-244 Kraków, Poland, chris@oa.uj.edu.pl, ORCID ID: <https://orcid.org/0000-0002-6280-2872>

Tomasz Kundera 

Astronomical Observatory of Jagiellonian University, ul. Orla 171, 30-244 Kraków, Poland,
chris@oa.uj.edu.pl, ORCID ID: <https://orcid.org/0000-0003-4593-1530>

Carole Roskowiński

Nicolaus Copernicus University in Toruń, ul. Gagarina 11, 87-100 Toruń, Poland,
carosko@astro.uni.torun.pl

Marta A. Dzielak

Nicolaus Copernicus Astronomical Center, Polish Academy of Sciences, ul. Bartycka 18,
00-716 Warszawa, Poland, mdzielak@camk.edu.pl

Received: 04.08.2017

Revised: 11.03.2019

Accepted: 11.03.2019